

REMARKS

Claims 32-35 and 37-56 are pending in the application. Claims have been rejected under 35 U.S.C. §103(a) as being deemed unpatentable over Beshai et al. (U.S. Patent No. 6,721,271), Massa (U.S. Patent No. 6,721,271, Edholm (U.S. Patent No. 6,721,271), Mauger et al. (U.S. Patent No. 6,917,586) and Wang (U.S. Patent No. 6,477,612.) Of the Claims, Claims 32, 40, 44, and 48 are independent. Claims have been amended to clarify the Applicants' invention. The application as amended and argued herein, is believed to overcome the rejections.

Regarding Claim Status

The Office indicates that Claims 32-56 are pending. However, Applicant notes that only Claims 32-35 and 37-56 are pending because Claim 36 was canceled in the amendment filed on May 29, 2007.

Regarding rejections under 35 U.S.C. §103(a)

Claims 32-35, 37-43 and 48-55 have been rejected under 35 U.S.C. §103(a) as being deemed unpatentable over Beshai et al. (U.S. Patent No. 6,721,271) in view of Massa et al. (U.S. Patent No. 6,658,469) and further in view of Edholm (U.S. Patent No. 6,600,721).

Claims 44-47 have been rejected under 35 U.S.C. §103(a) as being deemed unpatentable over Beshai et al. (U.S. Patent No. 6,721,271), in view of Massa et al. (U.S. Patent No. 6,658,469), in view of Edholm (U.S. Patent No. 6,600,721) and further in view of Mauger (U.S. Patent No. 6,917,586).

Claim 56 has been rejected under 35 U.S.C. §103(a) as being deemed unpatentable over Beshai et al. (U.S. Patent No. 6,721,271), in view of Massa et al. (U.S. Patent No. 6,658,469), in view of Edholm (U.S. Patent No. 6,600,721) in view of Mauger (U.S. Patent No. 6,917,586) and further in view of Wang (U.S. Patent No. 6,477,612.)

An embodiment of the Applicant's claimed invention improves efficiency of RDMA by coalescing (or merging) small RDMA requests into a single memory to memory transfer operation and dividing large RDMA requests into multiple operations to support hardware and

software limitations. (See, for example, Page 15, lines 8-15 of the Applicants' application as originally filed.)

In an embodiment of the Applicants' claimed invention shown in Fig. 7, if an amount of data located in a first memory buffer (70) in a local system associated with a remote direct memory access (RDMA) request (62) does not exceed a maximum transfer size for a single memory to memory transfer operation over a data network to a remote memory buffer in a remote system, the data is associated with a first transfer operation (84) and if the amount of data associated with the first transfer operation (84) has not reached the maximum transfer size, a portion of the data (72A) associated with a subsequent transfer operation (86) for the RDMA request (62) located in one or more portions of one or more other memory buffers (72) in the local system is associated with the first transfer operation (84) to the remote memory buffer. (See, Fig. 7 and Page 15, line 16 – Page 20, line 9 of the Applicants' application as originally filed.)

To establish a prima facie case for obviousness under 35 U.S.C. 103(a), (1) there must be some suggestion or motivation to combine reference teachings; (2) there must be a reasonable expectation of success; (3) the references when combined must teach or suggest all the claim limitations. For the reasons discussed below, it is respectfully submitted that the Office has not established a prima facie case under 35 U.S.C. 103(a) for claims 32-35 and 37-55 and that therefore, claims 32-35 and 37-55 are allowable.

The references when combined do not teach or suggest all the claim limitations

Beshai has been cited for disclosing the Applicants' claimed "if an amount of data located in a first memory buffer in a local system...does not exceed a maximum transfer size for a single memory to memory...associating the data with a first transfer operation; and if the amount of data associated with the first transfer operation has not reached the maximum transfer capacity, associating data for the RDMA request located in one or more portions of other memory buffers with the first transfer operation". (See Office Action mailed August 6, 2007, Page 2 – Page 3.)

Beshai discusses processing of received variable-sized network packets by dividing the packets into equal sized packet segments and aggregating packet segments destined for a particular egress module into a parcel. "[P]arcel is formed from an ingress buffer only when

the number of segments in the buffer equals or exceeds the parcel size, q, or if the waiting time of the head-of-buffer segment has reached a predetermined threshold” (See, Beshai, col. 12, lines 34-38 (emphasis added).)

Beshai does not teach or suggest at least:

“associating the data with a first transfer operation; and
if the amount of data associated with the first transfer operation has not reached the maximum transfer size, associating a portion of the data associated with a subsequent transfer operation for the RDMA request located in one or more portions of one or more other memory buffers in the local system with the first transfer operation to the remote memory buffer.”

as claimed by the Applicants in amended Claim 32.

In contrast to the applicants’ claimed invention, Beshai is merely directed to managing packet throughput within a 3-stage switch, by dividing packets into predetermined equal size segments and aggregating packet segments for a same egress module destination. “Segments are grouped into logical buffers according to egress-module destination”. Beshai discusses coalescing packet segments destined to an egress module in a switch in a parcel “only when the number of packet segments in buffer equals or exceeds the parcel size.” (See, col. 12, lines 34 – lines 38.)

Beshai does not teach or suggest “associating a portion of the data associated with a subsequent transfer operation for the RDMA request located in one or more portions of one or more other memory buffers in the local system with the first transfer operation to the remote memory buffer”. In contrast, Beshai merely discusses transferring packets received by ingress modules to egress modules within a switch based on “the number of parcels that each ingress module 32 is permitted to transfer to each egress module 36 during a specified transfer allocation period”.

In contrast, data transferred within the switch from ingress modules to egress modules is based on packet segments of equal size. (See, Beshai, for example, Fig. 1, column 8, line 32 – column 9, line 9.) Beshai merely discusses aggregation of these equal size packet segments that are destined to egress from the same egress module into parcels with the number of packet segments transferred based on for example, “a nominal capacity of 100 packet segments per transfer allocation”. (See, col. 15, lines 9-14.)

Massa has been cited for disclosing “transfer operation to a remote memory in a remote system with other memory buffers in the local system”. However, Massa does not teach or suggest at least:

“associating the data with a first transfer operation; and
 if the amount of data associated with the first transfer operation has not reached the maximum transfer size, associating a portion of the data associated with a subsequent transfer operation for the RDMA request located in one or more portions of one or more other memory buffers in the local system with the first transfer operation to the remote memory buffer.”

as claimed by the Applicants in amended Claim 32.

Massa’s discussion of two methods of transferring data (RDMA for large data transfer and messages for small data transfer) teaches away from the Applicants’ claimed method which coalesces small data transfers to transfer the maximum RDMA data in a single transfer operation. In contrast, Massa teaches a system that transfers multiple messages to transmit data that is not large enough to be transferred using an RDMA transfer operation. “The switch 126 detects the data size and decides whether to use RDMA data transfer or messages to transfer the data to application 132”. (See, Massa col. 11, lines 63-67.) Massa does not teach or suggest how much data is transferred in an RDMA transfer operation. Massa merely indicates that “remote switch 126 transfers an amount of data equal to the size of the set of receiving buffers 134 from the set of transmission buffers 138 into the set of receiving buffers 134 using one or more RDMA write operations”. (See, Massa, col. 12, lines 48-52.) Thus, Massa merely describes a prior art RDMA system as discussed in the background of the Applicants’ specification as originally filed, that is, “RDMA request can transfer data from multiple local memory buffers into a single remote memory region”. (See, Applicants’ specification, Page 2, lines 10-12.) In the RDMA system discussed by Massa, there is no teaching or suggestion of reducing the number of RDMA transfer operations to transfer data from local buffer(s) to a remote buffer by “associating a portion of the data associated with a subsequent transfer operation for the RDMA request located in one or more portions of one or more other memory buffers in the local system with the first transfer operation to the remote memory buffer” as claimed by the Applicants in claim 32 (as amended) (emphasis added).

The Office has cited Edholm for disclosing “data is to be transferred in a single transfer operation to a remote memory buffer”. Edholm merely discusses bandwidth management by adjusting latency between packets in order to reduce data flow (decrease the transmission rate). “[T]he receiving device defines a predetermined bandwidth and maintains the transmission at the transmitting device to remain within the predetermined bandwidth thereby minimizing the complicated managing of buffers used by an on and off transmission of flow control signals”. (See Edholm, col. 3, lines 29-34.) The latency between packets is controlled by a control application. “When control application 332 determines that the bandwidth limitations of the transmission rate have not been exceeded, the control application releases the pointer to driver 316. When the bandwidth limitations have been exceeded, the control application waits before placing the data packet on the transmit queue for driver 316. Thus the bandwidth limitations determine the latency between data packets. Higher bandwidth limit results in higher data transmission rates thereby reducing the latencies between data packets.” (See, Edholm, col. 4, lines 9-25.)

Edholm merely discusses packet processing, that is, a method for transmitting packets over a network, between NIC cards 312 coupled to a network 308. There is no discussion of a “single transfer operation to a remote memory buffer”. The Applicants’ claimed invention relates to Remote Direct Memory Access (RDMA), that is, memory to memory transfer. One skilled in the art of RDMA would not look to bandwidth management of a network connection between an end node and a switch to increase throughput in an RDMA application.

Claims 33-35 and 37-39 are dependent claims that depend directly or indirectly on claim 32 which has already been shown to be non-obvious over the cited art.

Furthermore, Beshai does not teach or suggest a “descriptor” as claimed by the Applicants in dependent claim 34. In contrast, Belshai merely describes “core memories 66” (for storing packet segments for a particular egress module in a switch. (See, col. 10, lines 48-65.) Belshai’s “core memories” does not teach or suggest the Applicants’ claimed “descriptor”. (See, for example, Fig. 7, descriptor (80).) The Applicants’ descriptor specifies “the remote memory buffer to which the data is to be transferred and to indicate a portion of data remaining to be transferred for the RDMA request”.

Independent claims 40, 44 and 48 recite a like distinction and are thus patentably distinguished over the cited art. Claims 41-43 depend directly or indirectly on claim 40, claims

45-47 depend directly or indirectly on claim 44 and claims 49-55 depend directly or indirectly on claim 48 and are thus patentably distinguished over the cited references.

Beshai has been cited for its teaching of packet management within a switch. Mauger has been cited for its teaching of “a host fabric adaptor”. Edholm has been cited for its teaching of bandwidth management by adjusting latency between packets. Massa has been cited for its teaching of RDMA operations. Wang has been cited for its teaching of virtual memory management. One skilled in the art of RDMA operations would not look to packet management within a switch, host fabric adaptors, adjusting latency between packets or virtual memory management to perform “a single memory to memory transfer operation over a data network to a remote memory buffer in a remote system” by “associating the data with a first transfer operation; and if the amount of data associated with the first transfer operation has not reached the maximum transfer size, associating a portion of the data associated with a subsequent transfer operation for the RDMA request located in one or more portions of one or more other memory buffers in the local system with the first transfer operation to the remote memory buffer”.

Therefore, separately or in combination, Beshai, Edholm, Massa, Mauger and Wang do not teach or suggest the applicants’ claimed invention. Even if combined, the present invention as now claimed does not result as argued above.

Thus, applicants respectfully request that the rejection of Claims 32-35, 37-43 and 48-55 over Beshai et al. in view of Edholm, Claims 44-47 over Beshai, Edholm in view of Mauger and Claim 56 over Beshai in view of Massa, Edholm and Mauger and further in view of Wang be withdrawn.

Accordingly, the present invention as now claimed is patentably distinguished from the cited references. Removal of the rejections under 35 U.S.C. § 103(a) and acceptance of claims 32-35 and 37-56 is respectfully requested.

CONCLUSION

In view of the foregoing, it is submitted that all claims (claims 32-35 and 37-56) are in condition of allowance. The Examiner is respectfully requested to contact the undersigned by telephone if such contact would further the examination of the above-referenced application.

Should an extension of time be necessary to respond to the outstanding Office Action, applicants respectfully petition for an extension of time pursuant to 37 C.F.R. § 1.136(a). Please charge our Deposit Account No. 50-0221 to cover the fee for the extension.

Respectfully submitted,

Date: 10/31/2007

/Caroline M. Fleming/
Caroline M. Fleming
Reg. No. 45,566
Telephone No. (978) 553-7371